Polygenic adaptation after a sudden change in environment

Laura K. Hayward [a, 1] and Guy Sella [b, c, d, 1]

[a] Department of Mathematics, Columbia University, New York, NY 10027

[b] Department of Biological Sciences, Columbia University, New York, NY 10027

[c] Department of Systems Biology, Columbia University, New York, NY 10032

[d] Program for Mathematical Genomics, Columbia University, New York, NY 10032

[1] To whom correspondence should be addressed: lauhayward@gmail.com or gs2742@columbia.edu

Abstract

Polygenic adaptation in response to selection on quantitative traits is thought to be ubiquitous in humans and other species, yet this mode of adaptation remains poorly understood. We investigate the dynamics of this process, assuming that a sudden change in environment shifts the optimal value of a highly polygenic quantitative trait. We find that when the shift is not too large relative to the genetic variance in the trait and this variance arises from segregating loci with small to moderate effect sizes (defined in terms of the selection acting on them before the shift), the mean phenotype's approach to the new optimum is well approximated by a rapid exponential process first described by Lande (1976). In contrast, when the shift is larger or large effect loci contribute substantially to genetic variance, the initially rapid approach is succeeded by a much slower one. In either case, the underlying changes to allele frequencies exhibit different behaviors short and long-term. Over the short term, strong directional selection on the trait introduces small differences between the frequencies of minor alleles whose effects are aligned with the shift in optimum versus those with effects in the opposite direction. The phenotypic effects of these differences are dominated by contributions from alleles with moderate and large effects, and cumulatively, these effects push the mean phenotype close to the new optimum. Over the longer term, weak directional selection on the trait can amplify the expected frequency differences between opposite alleles; however, since the mean phenotype is close to the new optimum, alleles are mainly affected by stabilizing selection on the trait. Consequently, the frequency differences between opposite alleles translate

into small differences in their probabilities of fixation, and the short-term phenotypic contributions of large effect alleles are largely supplanted by contributions of fixed, moderate ones. This process takes on the order of $\sim 4N_e$ generations (where N_e is the effective population size), after which the steady state architecture of genetic variation around the new optimum is restored.

Introduction

Many traits under natural selection are quantitative, highly heritable, and genetically complex, meaning that they take on continuous values, that a substantial fraction of the population variation in their values arises from genetic differences among individuals, and that this variation arises from small contributions at many segregating loci. It therefore stands to reason that the responses to changing selective pressures often involve adaptive changes in such traits, accomplished through changes to allele frequencies at the many loci that affect them. In other words, we should expect polygenic adaptation in complex, quantitative traits to be ubiquitous. This view traces back to the dawn of population and quantitative genetics (1, 2) and is supported by many lines of evidence (3, 4).

Notably, it is supported by studies of the response to directional, artificial selection on many traits in plants and animals in agriculture and in artificial evolution experiments (3, 4). In these settings, selected traits typically exhibit amazingly rapid and sustained adaptive changes (5-7), which are readily explained by models in which the change is driven by small shifts in allele frequencies at numerous loci (5, 8), and inconsistent with models with few alleles of large effects (6, 9). The potential importance of polygenic adaptation has also been highlighted by more recent efforts to elucidate the genetic basis of adaptation in humans. In the first decade after genome-wide polymorphism datasets became available, this quest was largely predicated on the monogenic model of a hard selective sweep (10, 11), in which adaptation proceeds by the fixation of new or initially rare beneficial mutations of large effects (e.g., (12)). Subsequent analyses, however, echoed studies of artificial selection in indicating that hard sweeps were rare, at least over the past ~500,000

years of human evolution (13, 14). Yet humans plausibly adapted in myriad ways during this time period, and they definitely experienced substantial changes in selection pressures, notably during more recent expansions across the globe. These considerations refocused the quest for the genetic basis of human adaptation on polygenic adaptation (15, 16).

Findings from genome wide association studies (GWASs) in humans have been central to this research program. Statistical analyses of GWASs indicate that in humans, heritable variation in complex traits is highly polygenic (17-19). For example, for many traits, estimates of the heritability contributed by chromosomes are approximately proportional to their length (17), suggesting that the contributing variants are numerous and roughly uniformly distributed across the genome. Such findings reinforced the view that adaptive changes to quantitative traits are likely to often be highly polygenic, but also implied that their identification would be difficult, as the changes to allele frequencies at individual loci may be minute. To overcome this limitation, recent studies pooled signatures of frequency changes over the hundreds to thousands of alleles that were found to be associated with an increase (or decrease) in a given trait (20-26). Initial studies suggest that polygenic adaptation has affected multiple human traits, but these conclusions have been called into question with the realization that the results are highly sensitive to systematic biases in GWASs, most notably due to residual population structure (27, 28).

Given that polygenic adaptation is plausibly ubiquitous, yet likely hard to identify, there is a clear need for a deep understanding of its behavior in populations and footprints in data. To date, theoretical work has primarily focused on two scenarios. The first is motivated by the observed responses to sustained artificial selection, modeled either as truncation selection (29) or as stabilizing selection, with the optimal phenotype moving at a constant rate in a given direction (e.g., (30-34)). In natural populations, however, quantitative traits are unlikely to be subject to long-term continuous change in one direction. Instead, considerable evidence indicates that they are often subject to long-term stabilizing selection (3), with intermittent shifts of the optimum in different directions. The second scenario therefore assumes that a sudden change in the environment induces an instantaneous shift in the optimum of a trait under stabilizing selection (35-42). Although

more elaborate scenarios (where, for example, the optimum and/or strength of stabilizing selection vary frequently) are also possible, this simple scenario provides a sensible starting point for thinking about polygenic adaptation in nature, and is our focus here.

Although there has been considerable work on the adaptive response to an instantaneous change in optimal phenotype, our understanding of this process is still limited. Seminal work by Lande (35) focused on the phenotypic response to selection in the infinitesimal limit, in which genetic variation arises from infinitely many segregating loci with infinitesimal effect sizes (see below). In reality, the number of loci and their effects are obviously finite. We would like to understand how this assumption affects the phenotypic response. Also, with GWASs now enabling us, at least in principle, to learn about the genetic basis of the phenotypic response, we would like to understand the allelic dynamics that underlie it.

Several studies have tackled this problem using simulations (e.g., (38, 39)). Although illustrative of the dynamics, it is unclear how to generalize their results, given (necessarily) arbitrary choices about multiple parameters and the complexity of these dynamics. In turn, elegant analytical work by de Vladar & Barton (36) and extensions by Jain & Stephan (37, 43) afford a general understanding of the allelic dynamics in models with an infinite population size. These dynamics, however, are shaped by features of mutation-selection balance that are specific to infinite populations. Notably, they strongly depend on the frequency of the allele prior to the shift in optimum following deterministically from its effect size, and on the critical effect size at which this frequency transitions from being dominated by selection to being dominated by mutation. But in real (finite) populations (including humans), the frequencies of alleles whose selection effects are sufficiently small to be dominated by mutation will be shaped by genetic drift; more generally, variation in allele frequencies due to genetic drift will crucially affect the allelic response to selection (see below). Thus, we still lack a solid understanding of the allelic dynamic underlying polygenic adaptation in natural populations.

Here, we follow previous work in considering the phenotypic and allelic responses of highly polygenic traits after a sudden change in optimal phenotype. But we do so in finite populations and employ a combination of analytic and simulation approaches to characterize how the responses varies across a broad range of evolutionary parameters.

The model

We build upon the standard model for the evolution of a highly polygenic, quantitative trait subject to stabilizing selection (3, 44-49). An individual's phenotype is represented by the value of a continuous trait, which follows from its genotype by the standard additive model (50, 51). Namely, we assume that the number of genomic sites affecting the trait (i.e., the target size) is very large, $L \gg 1$, and that an individual's phenotype is given by

$$z = \sum_{l=1}^{L}(a_l + a_l') + \epsilon, \qquad (1)$$

where the first term is the genetic contribution, with a_l and a_l' denoting the phenotypic effects of the parents' alleles at site l, and $\epsilon \sim N(0, V_E)$ is the environmental contribution.

Stabilizing selection is introduced by assuming that fitness declines with distance from the optimal trait value positioned at the origin ($z = 0$). Specifically, we assume a Gaussian (absolute) fitness function:

$$\mathrm{W}(z) = \exp(-z^2/2V_S), \qquad (2)$$

where V_s^{-1} measures the strength of selection. The specific form of the fitness function is unlikely to affect our results under parameter ranges of interest (see below), however. Since the additive environmental contribution to the phenotype can be absorbed into V_S (by replacing it by $V_S' = V_S + V_E$; e.g., (46, 52)), we consider only the genetic contribution.

The population dynamics follow the standard model of a diploid, panmictic population of constant size N, with non-overlapping generations. In each generation, parents are randomly chosen to reproduce with probabilities proportional to their fitness (i.e., Wright-Fisher sampling with fertility selection), followed by mutation, free recombination (i.e., no linkage) and Mendelian segregation. We assume that the mutational input per site per

generation is sufficiently small such that segregating sites are rarely more than bi-allelic (i.e., that $\theta = 4Nu \ll 1$, where u is the mutation rate per site per generation). We therefore employ the infinite sites approximation, such that the number of mutations per gamete, per generation follows a Poisson distribution with mean $U = Lu$. The effect sizes of mutations, a, are drawn from a symmetric distribution, i.e., with equal probability of increasing or decreasing the trait value; further assumptions about this distribution are specified below. Table S1 provides a summary of our notation.

Evolutionary scenario and parameter ranges. We assume that at the outset–before the shift in optimal phenotype–the population has attained mutation-selection-drift balance. We follow previous work on this balance in making several plausible assumptions about parameter ranges (e.g., (49)). First, we assume that the per generation, population scaled mutational input is sufficiently large for variation in the trait to be highly polygenic (specifically, that $\sqrt{2NU} \gg 1$). Second, we assume that the expected number of mutations affecting the trait per generation, per gamete, is substantially smaller than 1 ($U = Lu \ll 1$). For this assumption to be violated, a trait would have to be extremely polygenic, e.g., in humans, the mutational target size, L, would have to exceed ~5 Mb (assuming that $u \approx 2 \cdot 10^{-8}$ per bp per generation, including all types of mutations); while we believe that our results would still hold qualitatively when $U > 1$, we leave the investigation of this case for future studies. Third, we make the standard assumption that selection coefficients of all alleles satisfy $s \ll 1$, which implies that the equilibrium selection coefficient $s_e = a^2/V_S \ll 1$ (see below and (2, 46, 53)). Fourth, we assume that a substantial proportion of mutations are not effectively neutral, i.e., have $S = 2Ns_e \gtrsim 1$. This assumption is supported by empirically based estimates of persistence time for a variety of traits and taxa (3, 4) and by inferences based on human GWASs (49, 54), indicating that quantitative genetic variance is not predominantly neutral. Under these assumptions, the phenotypic distribution at mutation-selection-drift balance is symmetric and tightly centered on the optimal phenotype (Fig. 1). Specifically, the mean phenotype exhibits tiny, rapid fluctuations around the optimal phenotype with variance $\delta^2 = V_S/2N$ (49); the phenotypic standard deviation is considerably greater than these fluctuations, i.e., $\sqrt{V_A} \gg \delta$ (SI Section 2); but

the phenotypic variance is much smaller than the curvature of the fitness function, i.e., $V_S \gg V_A$ (49).

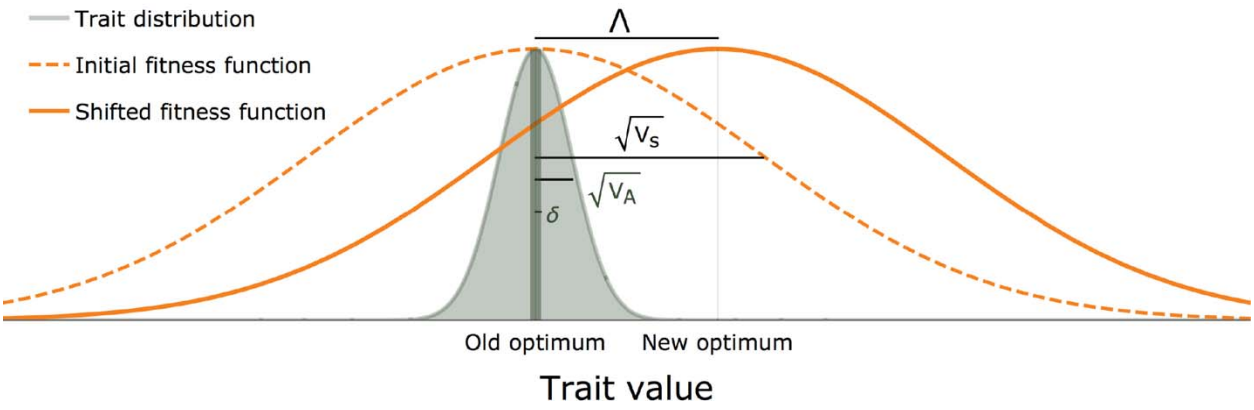

Figure 1. The evolutionary scenario. Before the shift in optimum, phenotypes are distributed symmetrically, with a mean that is very close to the old optimum and a variance that is much smaller than the curvature of the fitness function ($V_A \ll V_S$). We consider the response to an instantaneous shift in optimum, for the case where the magnitude of the shift is smaller than the width of the fitness function ($\Lambda < \sqrt{V_S}$). See text for further details.

We consider the response to an instantaneous shift of Λ in optimal phenotype at time t=0 (Fig. 1). We assume that the shift in optimum is greater than the equilibrium fluctuations in mean phenotype, i.e., that $\Lambda > \delta$. We further assume that $\Lambda < \sqrt{V_S}$ and $a \ll \sqrt{V_S}$. The latter requirements ensure that the maximal directional selection coefficients of alleles, which are attained immediately after the shift, satisfy $s_d = 2\,\Lambda \cdot a/V_S \ll 1$ (see below and (2, 55)). The requirement that $a \ll \sqrt{V_S}$ is not particularly restrictive, as it allows for the selection coefficients of alleles at equilibrium, $s_e = a^2/V_S$, to be as large as 1%. Neither is the assumption that $\Lambda < \sqrt{V_S}$, because, given that the genetic variance before the shift satisfies $V_A(0) \ll V_S$, it allows for shifts of several equilibrium phenotypic standard deviations (SI Section 2 and Fig. S1).

Choice of units. Our analysis allows us to choose the units in which we measure the trait. When we study the allelic response, we use units based on the dynamics at mutation-selection drift balance (before the shift in optimum). The population-scaled selection

coefficient at steady state is $S = 2Ns_e = 2N\tilde{a}^2/\tilde{V}_S = \tilde{a}^2/\delta^2$; here, the effect size $\tilde{a}$ is measured in arbitrary units, and $\tilde{V}_S$ is measured in these units-squared, such that the scaled selection coefficient has no units. We will measure the trait in units of δ. In these units, the effect size $a \equiv \tilde{a}/\delta = \sqrt{S}$, the stabilizing selection parameter $V_S \equiv \tilde{V}_S/\delta^2 = 2N$, and an allele's contribution to variance is $v^*(S, x) = 2Sx(1-x)$ (and has units of δ^2). We also measure the distance between the mean and optimal phenotype, D, and shift in optimum, Λ, in units of δ. Stating our results in these terms makes their form invariant with respect to changes in the population size, N, and the strength of stabilizing selection, V_S^{-1}.

Simulations and resources. We compare our analytical results to three layers of simulations (see SI Section 4 for further detail). The first realizes the full model described above, and is run with a burn in period of $10N$ generations to attain steady state before the shift in optimum. The second traces all alleles (AA) rather than individuals. It assumes linkage equilibrium (rather than free recombination), and changes to allele frequencies every generation are modeled according to the diffusion approximation detailed below. It is also run with a burn-in period of $10N$ generations before the shift. The third kind of simulation traces the dynamic of one allele at the time (OA). To that end: i) we sample initial minor allele frequencies from the closed form, equilibrium distributions ((49) and SI Section 4), using importance sampling based on the density of variance contributed by different minor allele frequencies (SI Section 4); and ii) the change in the population's mean phenotype over time, on which the allelic dynamics depend, is given as input, based on either an analytical approximation (see below) or on an average over simulations of the second layer. The last two layers allow for greater computational tractability, and their results are validated against those from the first layer (Fig. S2). Documented code for simulations, numerical analysis, and graphs can be found at https://github.com/sellalab/PolygenicAdaptation.

Results

Phenotypic response. We first consider how the population's mean phenotype approaches the new optimum. In SI Section 2, we express the mean distance from the new optimum, $D(t)$, as a sum over allelic contributions. We show that under our assumptions, the expected, per generation change in this distance is well approximated by

$$E\big(\Delta D(t)\big) \cong -(V_A(t)/V_S) \cdot D(t) + \mu_3(t)/(2V_S), \qquad (2)$$

where $V_A(t)$ and $\mu_3(t)$ denote the 2nd and 3rd central moments of the phenotypic distribution. Similar expressions were derived by Barton and Turelli (55) under the rare-alleles approximation and by Bürger (56) under the assumption of a parabolic fitness function.

We rely on Eq. 2 to describe the phenotypic response to selection. This response takes a simple form in the infinitesimal limit, in which genetic variation at equilibrium arises from infinitely many segregating alleles with infinitesimal effect sizes (57-59). In this limit, the change in mean phenotype is achieved by infinitesimal changes to allele frequencies at infinitely many loci; and as a result, the phenotypic distribution remains normal and the phenotypic variance remains constant. Under these assumptions, Eq. 2 reduces to

$$E\big(\Delta D(t)\big) = -(V_A(0)/V_S) \cdot D(t), \qquad (3)$$

which (in continuous time) is solved by

$$D_L(t) = \Lambda \cdot \exp(-(V_A(0)/V_S) \cdot t). \qquad (4)$$

This solution was first derived by Lande (35), and we refer to it henceforth as Lande's solution or approximation. When genetic variance is dominated by loci with small and intermediate effect sizes (as defined below) and the shift in optimum is not too large relative to the phenotypic standard deviation, changes to the 2nd and 3rd moments of the phenotypic distribution are small and the expected phenotypic response is well approximated by Lande's solution (Figs. 2A and S3).

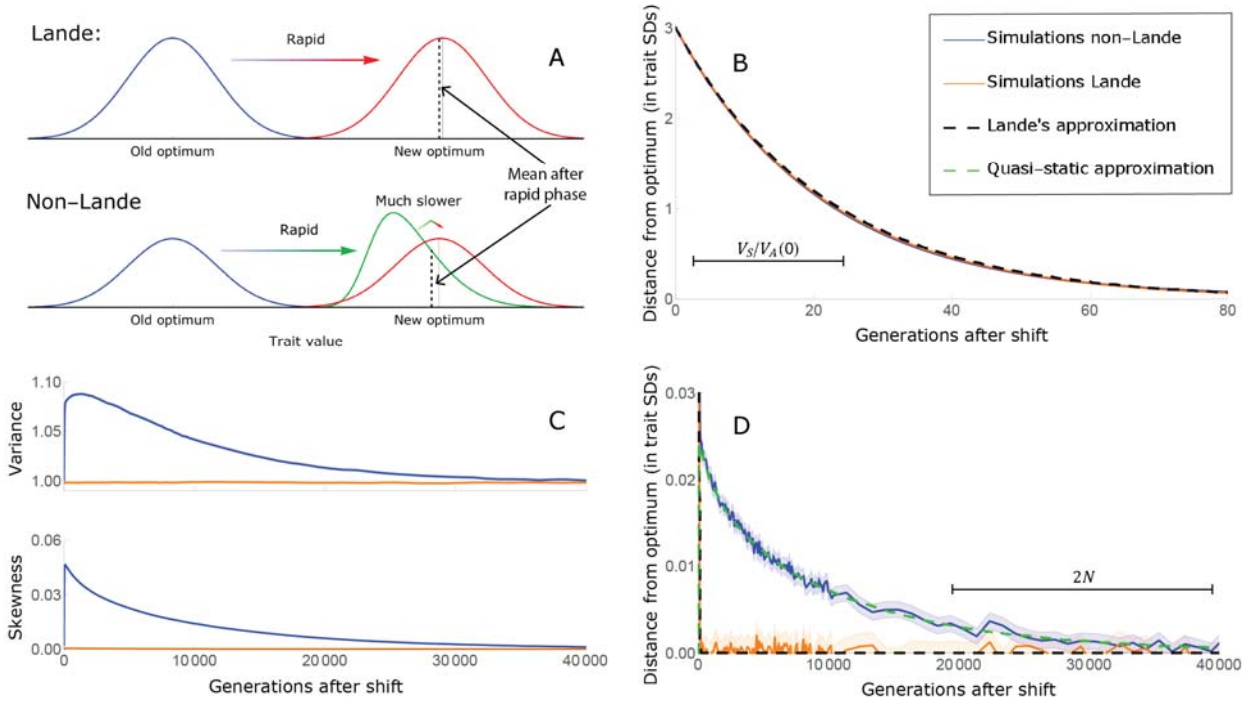

Figure 2. The phenotypic response to a shift in optimal phenotype. A) Cartoon of the two kinds of phenotypic response: i) the Lande approximation, in which the mean approaches the new optimum exponentially with time and the phenotypic distribution maintains its shape; ii) substantial deviations from Lande's approximation, in which the mean at first approaches the new optimum rapidly but in so doing, the phenotypic distribution becomes skewed, after which the mean's approach slows down to a rate that is dictated by the decay of the 3rd moment. B) In both the Lande and non-Lande cases, the mean phenotype initially approaches the new optimum rapidly. C) In the non-Lande case, the phenotypic variance and skewness increase during the rapid phase and then take a very long time to decay to their values at steady state. D) Over the longer term, the approach to the optimum in the non-Lande case almost grids to a halt, where its rate can be described by the quasi-static approximation (Eq. 5). The simulation results in B-D were averaged over 5000 runs of our allelic simulations (AA) (see Model section and SI Section 4), with $N = 10^4$, $\Lambda = 4 \cdot \sqrt{V_A(0)}$, and $\sqrt{V_A(0)} = 30 \cdot \delta$. In the Lande case, effect sizes were Gamma distributed with $E(S) = V(S) = 1$ ($S \sim \Gamma(1,1)$) and $U \cong 0.03$ (to match $V_A(0)$ in both cases); in the non-Lande case, effect sizes were Gamma distributed effect sizes with $E(S) = 24$ and $V(S) = 36$ ($S \sim \Gamma(6,4)$) and $U = 0.01$.

When alleles with large effects contribute markedly to genetic variance, or when the shift in optimum is large relative to the phenotypic standard deviation, changes to the 2nd and 3rd moments of the phenotypic distribution become more substantial (Fig. 2C; (55)). To see

why, consider a pair of minor alleles with the same effect size and initial frequency, where the effect of one is aligned with the shift and the effect of the other opposes it. After the shift, directional selection increases the frequency of the aligned alleles relative to that of the opposing one. The frequency increase of the aligned allele increases variance more than the frequency decrease of the opposing allele decreases it, resulting in a net increase to variance (Fig. 2C; (36, 43, 55)). The relative changes in frequency and thus the net increase in variance are greater for alleles with larger effects. Next consider the 3rd moment. At steady state, the contribution of alleles with opposing effects to the 3rd moment cancel out. After the shift, the frequency increase of aligned alleles relative to opposing ones introduces a non-zero 3rd moment (Fig. 2C). Large effect alleles contribute substantially more to this 3rd moment, both because their individual, steady state contribution to the 3rd moment is greater to begin with (SI Section 2) and because they exhibit larger relative changes in frequency after the shift (see below). By the same token, larger shifts result in stronger directional selection and greater relative frequency differences between alleles with opposing effects, and thus in greater increases to variance and a greater skew of the phenotypic distribution.

The increase in 2nd and 3rd moments after the shift result in a phenotypic dynamic with two distinct phases. First, immediately after the shift, the mean phenotype rapidly approaches the new optimum, akin to the exponential approach in Lande's approximation. In this case, however, genetic variance increases and thus the exponential rate of approach may increase, making the expected approach even faster (Eq. 2). Shortly thereafter, when the mean phenotype nears the optimum, the decreasing distance and increasing 3rd moment reach the point at which

$$D(t) \approx \mu_3(t)/(2V_A(t)). \quad (5)$$

The two terms on the right-hand side of Eq. 2 then approximately cancel out, and the dynamic enters a second, prolonged phase, in which the approach to the optimum nearly grinds to a halt (Fig. 2D). During this phase, the expected change in mean phenotype can be described in terms of a quasi-static approximation given by Eq. 5 (Fig. 2D). The rate of approaching the optimum is then largely determined by the rate at which the 3rd moment decays. This roughly corresponds to the rate at which the allele frequency distribution

equilibrates and steady state around the new optimum is restored (see section on "Equilibration time").

Allelic dynamics. We now turn to the allelic dynamics that underlie the phenotypic response. These dynamics can be described in terms of the first two moments of change in frequency in a single generation (60, 61). For an allele with effect size a and frequency x, we calculate the moments by averaging the fitness of the three genotypes over genetic backgrounds (SI Section 1). Under our assumptions, the moments are well approximated by

$$E(\Delta x) \cong (a \cdot D(t)/V_S) \cdot x(1-x) - (a^2/V_S) \cdot x(1-x)(\tfrac{1}{2} - x) \qquad (6)$$

and

$$V(\Delta x) \cong x(1-x)/2N, \qquad (7)$$

which is the standard drift term. Similar expressions for the first moment trace back to Wright (44) and have been used previously to study the response to selection on quantitative traits (36, 56, 62, 63).

The two terms in the first moment reflect different modes of selection: directional and stabilizing, respectively. The first term arises from directional selection on the trait and takes a semi-dominant form with selection coefficient $s_d = 2a \cdot D(t)/V_s$. Its effect is to increase the frequency of alleles whose effects are aligned with the shift (and vice versa) and its strength weakens as the distance to the new optimum, D, decreases. The second term arises from stabilizing selection on the trait and takes an under-dominant form with selection coefficient $s_e = a^2/V_S$. Its effect is to decrease an allele's contribution to phenotypic variance, $2a^2x(1-x)$, by reducing minor allele frequency (MAF); it becomes weaker as the MAF approaches ½.

The relative importance of the two modes of selection varies as the mean distance to the new optimum, D, decreases. We therefore divide the allelic response into two phases: a *rapid phase*, immediately after the shift, in which the mean distance to the new optimum is substantial and changes rapidly, and a subsequent, prolonged *equilibration phase*, in which

the mean distance is small and changes slowly (SI Section 3; (43)). We delimit the rapid phase by the time at which the distance $D(t_1)$ first equals $\delta = \sqrt{V_s/2N}$ or $\mu_3(t_1)/V_A(t_1)$. When Lande's approximation applies, the distance $D(t_1) = \delta$, where δ equals the standard deviation of the distance from the optimum at steady state (when $D = 0$). When deviations from Lande's approximation are substantial, the distance $D(t_1) = \mu_3(t_1)/V_A(t_1) > \delta$, where $\mu_3(t_1)/V_A(t_1)$ is the distance at which the expected change in D is halved due to the increase in the 2nd and 3rd moments of the phenotypic distribution (Eq. 2). In the latter case, the distance D during the equilibration phase will be larger, but under plausible parameter values it will be small in either case (SI Section 2). While the definition of the end of the rapid phase is somewhat arbitrary, our analysis is insensitive to the choice.

The change in mean phenotype during the rapid phase is driven by the differential effect of directional selection on minor alleles whose effects are aligned and opposed to the shift in optimum (Fig. 3). Considering a pair of minor alleles with opposing effects and the same initial frequency, selection increases the frequency of the aligned allele relative to the opposing one. By the end of the rapid phase, the frequency differences across *all* aligned and opposing alleles drive the mean phenotype close to the new optimum (Fig. 2A). Deviations from Lande's approximation manifest as prolonged, weak directional selection during the equilibration phase, which further increases the expected frequency difference between aligned and opposing alleles. However, given that we are considering a highly polygenic trait, the expected frequency difference between a pair of opposing alleles will be small. This small difference causes aligned alleles to have a slightly greater probability of eventually fixing during the equilibration phase (Fig. 3). Over a period on the order of $4N$ generations (see below), the frequency differences between aligned and opposing alleles are replaced by a slight excess of fixed differences between them, and the steady state genetic architecture is restored around the new optimum. In the following sections, we describe these processes quantitatively. Specifically, we ask how the relative contribution of alleles to phenotypic change during the two phases depends on their effect size and initial frequency.

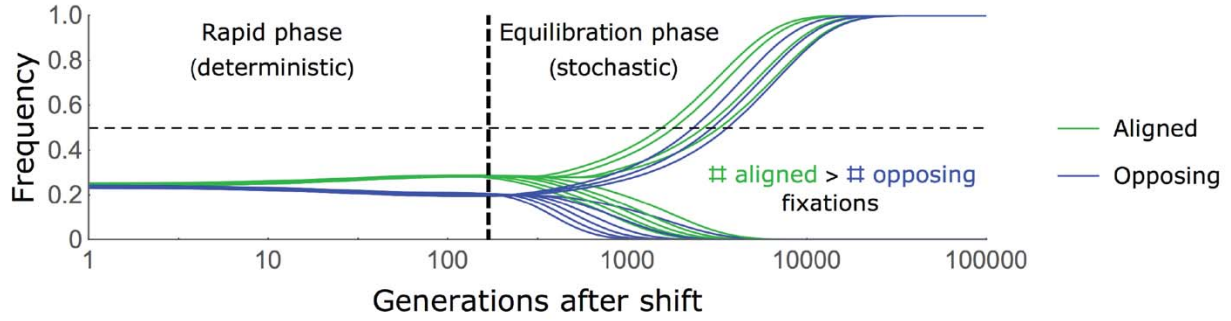

Figure 3. A cartoon of the allelic dynamic. We divide the allelic dynamics into rapid and equilibration phases, based on the rate of phenotypic change, and consider the trajectories of alleles with opposing effects of the same size, which start at the same initial minor frequency. During the rapid phase, alleles whose effects align with the shift slightly increase in frequency relative to those with opposing effects. During the equilibration phase, this frequency difference can increase further and eventually leads aligned alleles to fix with slightly greater probabilities than opposing ones.

The allelic response in the rapid phase. We can describe changes to allele frequencies during the rapid phase with a simple deterministic approximation. The duration of the rapid phase is much shorter than the time scale over which genetic drift has a substantial effect ($t_1 \sim 1/U \ll 2N$ generations; SI Section 3.2), allowing us to rely only on the first moment of change in allele frequency (Eq. 6). Additionally, deviations of the distance $D(t)$ from Lande's approximation during this phase have negligible effects (Fig. 2B and SI Section 2), allowing us to assume that $D(t) = D_L(t)$ (Eq. 4). Lastly, when relative frequency changes are small, we can substitute the frequency in the first moment by its initial value. With these simplifications, we can integrate the first moment over time to obtain an explicit approximation for frequency changes.

Consider a pair of minor alleles with opposing effects of size a and initial frequency x_0 before the shift in optimum. Using our simple approximation, we find that the frequency difference between them at the end of the rapid phase is

$$\Delta x^*_{t_1}(S, x_0) = x^+_{t_1}(S, x_0) - x^-_{t_1}(S, x_0) \approx 2(a/V_S) \cdot x_0(1 - x_0) \int_0^{t_1} D_L(\mathrm{t}) dt$$
$$= \left(\Lambda - D_L(t_1)\right) \cdot 2\, a x_0(1 - x_0)/V_A(0). \qquad (8)$$

Similarly, the contribution of the pair to the change in mean phenotype is

$$\Delta z^*_{t_1}(a, x_0) = 2a \cdot \Delta x^*_{t_1}(a, x_0) \approx \left(\Lambda - D_L(t_1)\right) \cdot 2\, v^*(a, x_0)/V_A(0), \qquad (9)$$

where $v^*(a, x_0) = 2a^2 x_0(1 - x_0)$ is the contribution to variance of an allele with effect size a and frequency x_0. Thus, the pair's contribution to phenotypic change is proportional to its contribution to phenotypic variance before the shift in optimum.

The expected relative contribution of *all* alleles with a given effect size and initial frequency is therefore proportional to their expected initial, steady state contribution to phenotypic variance. We will focus on the contribution per unit mutational input of alleles with a given effect size. To this end, we measure the trait value in units of $\delta = \sqrt{V_S/2N}$ and express allelic effect sizes in terms of the scaled selection coefficients at steady state (when $D = 0$), $S = 2Ns_e = a^2$ in these units (see Model section). Expressing our results in this form makes them invariant with respect to changing the population size, N, stabilizing selection parameter, V_S, mutational input per generation, $2NU$, and distribution of effect sizes, $g(S)$. In these terms, the expected contribution of alleles with given effect size and initial MAF to phenotypic change is

$$\Delta z_{t_1}(S, x_0) \approx \left(\Lambda - D_L(t_1)\right) \cdot v(S, x_0)/V_A(0), \qquad (10)$$

and the marginal contribution of alleles with a given effect size is

$$\Delta z_{t_1}(S) = \int_0^{1/2} \Delta z_{t_1}(S, x)dx \approx \left(\Lambda - D_L(t_1)\right) \cdot v(S)/V_A(0), \qquad (11)$$

where $v(S, x_0) \approx 4Se^{-Sx_0(1-x_0)}$ and

$v(S) \approx 4S \cdot \int_0^{1/2} \exp\left(-Sx(1-x)\right) dx = 4\sqrt{S}\,\mathrm{DawsonF}[\sqrt{S}/2]$ are the corresponding densities of variance per unit mutational input at steady state (SI Section 3.1). The absolute expected contributions follow from multiplying these expressions by the mutational input per generation, $2NU \cdot g(S)$. Specifically, as we would expect, the total change in mean phenotype during the rapid phase is $\Delta z_{t_1} = 2NU \cdot \int \Delta z_{t_1}(S) \cdot g(S)dS = \Lambda - D_L(t_1)$ (as $V_A(0) = 2NU \cdot \int v(S) \cdot g(S)ds$).

The relative contribution of alleles with given effect size and initial MAF to phenotypic change follows from their expected contribution to variance at steady state (Eqs. 10 and 11,

and Fig. 4). The properties of $v(S)$ imply that (Fig. 4A): i) the relative contribution of alleles with small effect sizes ($S \ll 1$) scale linearly with S ($v(S) \cong 2S$, measured in units of δ^2); ii) the contribution of alleles with moderate and large effect sizes (roughly $S = 2Na^2 > 3$) are on the same order ($v(S)\sim 4$) and much greater than for small effect ones; and iii) the contribution is maximized for $S \approx 10$ ($v(10) \approx 5.2$) (see (49) for intuition about these properties). While large and moderate effect alleles make similar contributions to phenotypic change, the MAFs of large effect alleles before the shift are much lower than the MAFs of moderate ones (Fig. 4B), because they are subject to stronger stabilizing selection. The expected frequency difference between pairs of opposing alleles is greatest for moderate effect sizes (Fig. 4C), because it is proportional to $E(2ax_0(1-x_0)) = \left(1/\sqrt{S}\right) \cdot v(S)$ (Eq. 8) and $v(S)$ is similar for moderate and large effect sizes. Additional properties of the allelic response during the rapid phase are presented in Fig. S4.

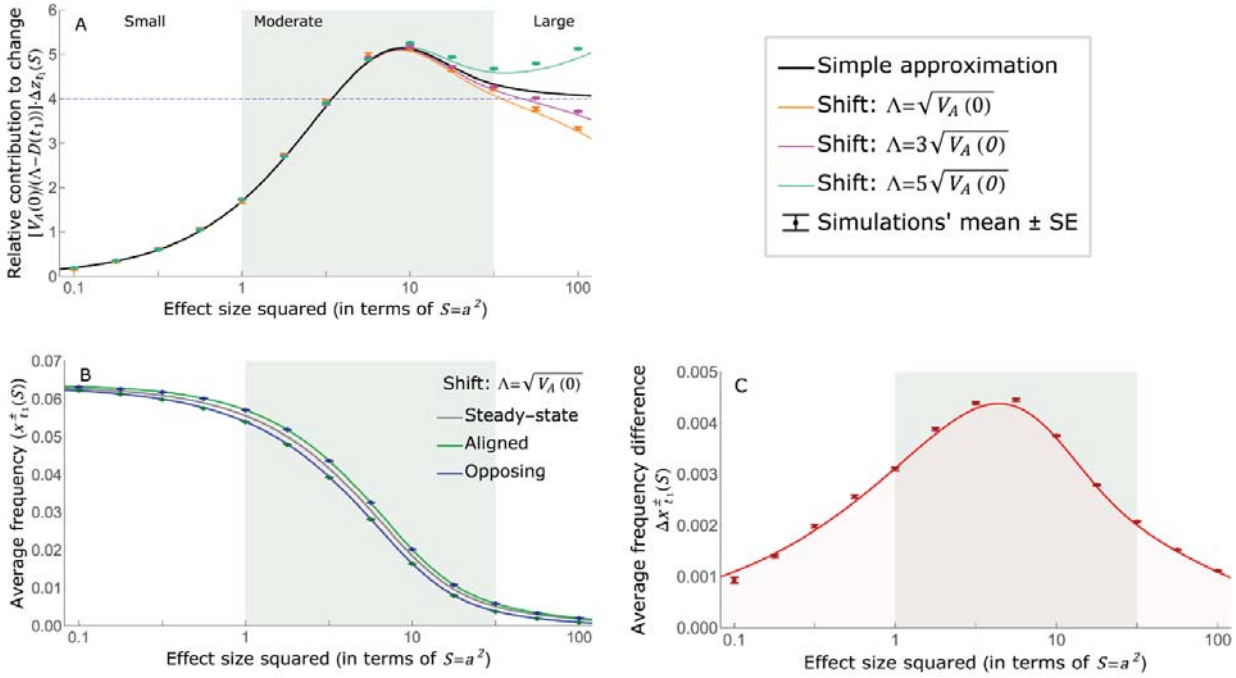

Figure 4. The allelic response during the rapid phase. A) Alleles with moderate and large effects make the greatest contribution to phenotypic change (per unit mutational input). The results of our simple approximation are compared with a more accurate one (see SI Section 3.2) and with simulations. B) The average MAF of aligned and opposing alleles at the end of the rapid phase decreases with effect size. C) The expected frequency difference

between pairs of opposing alleles is greatest for moderate effect sizes. Simulation results for each point were averaged over 10^5 runs of our individual allele (OA) simulation, assuming $N = 10^4$, Lande's approximation with $\sqrt{V_A}$=25δ, and the shifts specified in the legend box.

When the shift in optimum is large relative to the phenotypic standard deviation, our simple approximation becomes less accurate (Fig. 4A). This is most pronounced for alleles with large effects, which experience greatest change relative to their MAF before the shift (as a result of which, substituting the initial frequency into the 1st moment entails greater inaccuracy). These cases can be accurately described using more elaborate approximations (Fig. 4A and SI Section 3.2). Importantly, the qualitative behaviors we outlined are also seen.

The allelic response in the equilibration phase. Over the long run, the small frequency differences between opposite alleles that accrued in the rapid phase translate into small differences in their fixation probabilities (Fig. 3). Weak directional selection during the equilibration phase amplifies these differences in fixation probabilities. We begin by treating fixations under the assumption that Lande's approximation works well, which implies that this amplification is negligible, after which we turn to the general case. When Lande's approximation applies, we can approximate the fixation probability of an allele with effect size a and initial frequency x_0 in two steps. First, we use our deterministic approximation to calculate its frequency at the end of the rapid phase, x_1. Second, because directional selection has a negligible effect during the equilibration phase, we assume that $D = 0$ and rely on the diffusion approximation to derive the fixation probability $\pi(S, x_1)$ (SI Section 3.2).

Consider a pair of opposite minor alleles, with effect size a (and corresponding S) and initial frequency x_0, which reach frequencies x_1^+ and x_1^- by the end of the rapid phase. If we assume that the frequency changes during the rapid phase are small, we can approximate the pair's expected long-term, fixed contribution to phenotypic change by

$$\Delta z_\infty^*(S, x_0) \approx 2a\left(\pi(S, x_1^+) - \pi(S, x_1^-)\right)$$

$$\approx 2a\frac{\partial\pi}{\partial x}(S,x_0)\cdot\Delta x^*_{t_1}(a,x_0)=\frac{\partial\pi}{\partial x}(S,x_0)\cdot\Delta z^*_{t_1}(a,x_0). \quad (12)$$

In SI Section 3.2 we show that

$$\frac{\partial\pi}{\partial x}(S,x_0)=2f(S)/v(S,x_0),$$

where $f(S)\equiv 2S^{3/2}\cdot e^{-S/4}/(\sqrt{\pi}\cdot \mathrm{Erf}(\sqrt{S}/2))$ and $v(S,x_0)$ is the steady state density of variance per unit mutational input. From Eqs. 10-12, we find that the expected fixed contribution per unit mutational input of alleles with given effect size and initial MAF is

$$\Delta z_\infty(S,x_0)\approx\frac{\partial\pi}{\partial x}(S,x_0)\cdot\Delta z_{t_1}(S,x_0)\approx 2\left(\Lambda-D_L(t_1)\right)f(S)/V_A(0)$$

(note that this expression does not depend on the initial frequency). Further assuming that weak directional selection during the equilibration phase has the same proportional effects on expected frequency differences as directional selection during the rapid phase, and thus replacing $\Lambda-D_L(t_1)$ by Λ, we find that

$$\Delta z_\infty(S,x_0)\approx 2\Lambda f(S)/V_A(0). \quad (13)$$

The expected marginal contribution of alleles with a given effect size follows and is

$$\Delta z_\infty(S)=\int_0^{1/2}\Delta z_\infty(S,x)dx\approx\Lambda f(S)/V_A(0). \quad (14)$$

When large effect alleles contribute markedly to genetic variance before the shift, however, this approximation substantially underestimates the long-term change in mean phenotype. To see why, note that $f(S)<v(S)$ for any S, but the difference becomes large when $S\gtrsim 4$ (Fig. 5A and SI Section 3.2). If the bulk of variance before the shift arises from alleles with $S<4$ then $V_A(0)\approx 2NU\cdot\int_0^4 v(S)\cdot g(S)dS\approx 2NU\cdot\int_0^4 f(S)\cdot g(S)dS$. Our approximation for the long-term change in mean phenotype then satisfies

$$\Delta z_\infty=2NU\cdot\int\Delta z_\infty(S)\cdot g(S)dS\approx\Lambda\cdot(2NU\cdot\int f(S)\cdot g(S)dS)/V_A(0)\approx\Lambda.$$

Hence, our approximation for the change in mean is only slightly smaller than the shift in optimum (Fig. 5A). However, when large effect alleles contribute markedly to genetic variance, the same reasoning implies that our approximation for the change in mean is substantially smaller than the shift, and is thus a substantial underestimate. In this case, deviations from Lande's approximation are substantial: the prevalence of large effect alleles leads to a quasi-static decay of the distance D during the equilibration phase (Fig. 2D and section on "phenotypic response"), and the prolonged, weak directional selection that

results markedly amplifies the difference in fixation probabilities between opposite alleles. This amplification, which is not accounted for in our simple approximation, ensures that the long-term change in mean phenotype equals the shift in optimum (Fig. S5 and SI Section 3.3).

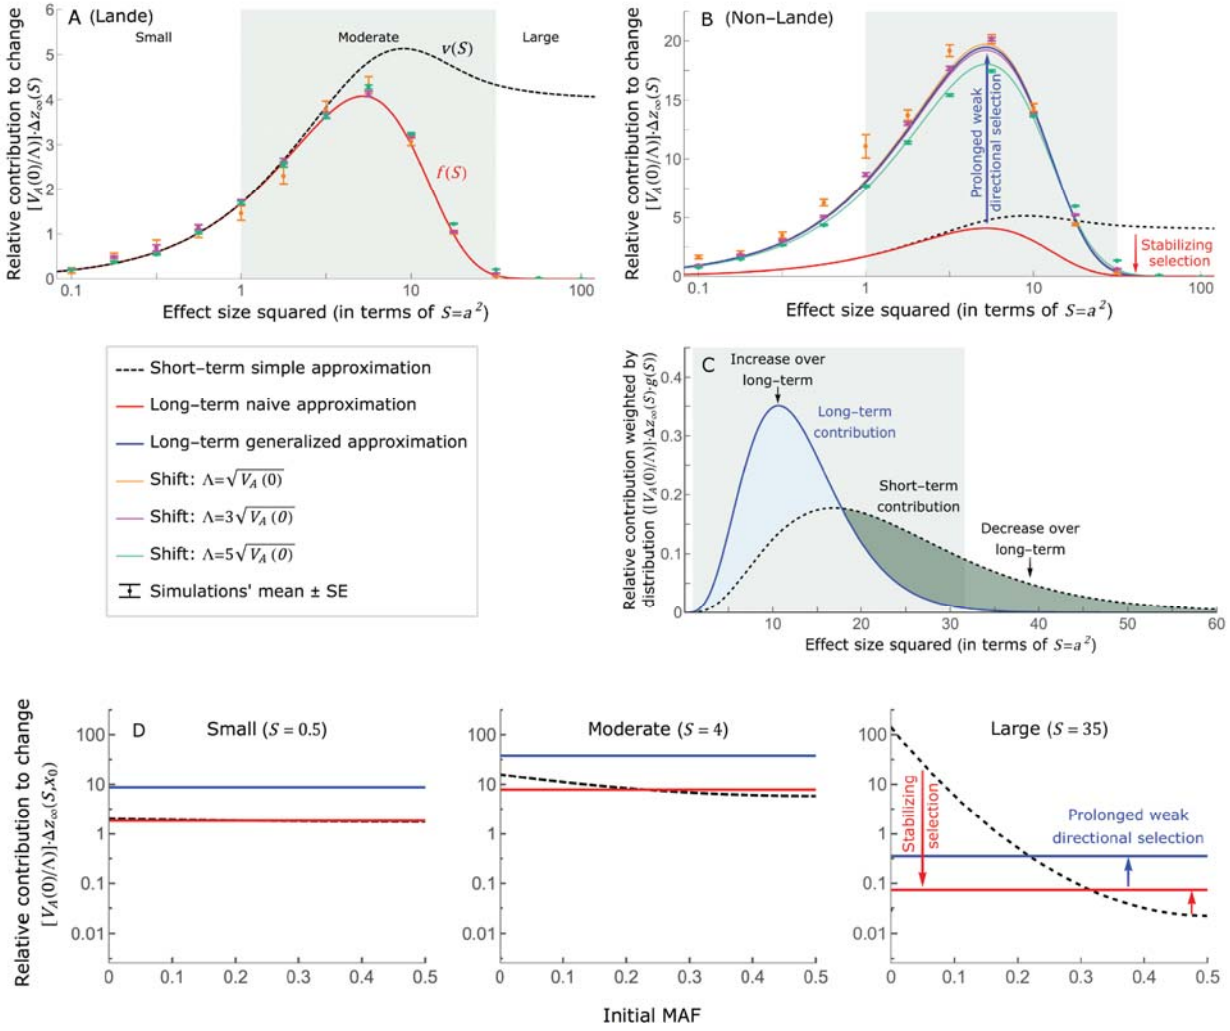

Figure 5. Contrasting the short and long-term allelic contributions to phenotypic adaptation reveals massive turnover in the genetic basis of adaptation during the equilibration phase. A) Comparison of the relative short and long-term contributions of alleles as a function of effect size, assuming Lande's approximation (Eqs. 11 and 14). Simulation results for each point were averaged over 10^5 runs of our individual allele simulation (OA), assuming Lande's approximation with $\sqrt{V_A}=25\delta$, $N = 10^4$, and the shifts specified in the legend box. While our simple approximation for the fixed contributions performs well in these examples, the theoretical curve for the long-term contribution (corresponding to $f(S)$) is always below the curve for the short-term contribution

(corresponding to $v(S)$), and the difference becomes substantial for $S \gtrsim 4$. In other words, the approximation underestimates the fixed contribution when large effect alleles contribute markedly to genetic variance. B) Comparison of the relative short and long-term contributions as a function of effect size, using our generalized approximation for fixed contributions (Eq. 16). We assume the effect size distribution corresponding to the non-Lande case described in Fig. 2, which has an amplification factor of $C \approx 4.77$. Long-term stabilizing selection diminishes the contribution of alleles with large effects and amplifies that of alleles with small and moderate effects (red arrow); long-term, weak directional selection amplifies the contribution of alleles with small and moderate effects (blue arrow). Simulation parameters also correspond to the non-Lande case in Fig. 2, with the shift sizes indicated in the legend. Simulation results for each point were averaged over 10^5 runs of our individual allele simulation (OA), where $D(t)$ was averaged over 10^3 allelic simulations (AA). C) The same results presented in B, with the short- and long-term relative contributions weighted by the mutational distribution of effect sizes ($S \sim \Gamma(6,4)$) and squared effect sizes (on the x-axis) shown on a linear rather than a log scale. The shaded blue area corresponds to the increase in the contribution of moderate effect alleles over the long-term and the shaded dark gray area corresponds to the decrease in the contribution of large effect alleles over the long-term; their equal size illustrates how the short-term contribution of large effect alleles is supplanted by longer-term fixations of moderate effect alleles. D) Comparison of the relative short and long-term contributions as a function of initial MAF, based on our approximations (Eqs. 10 and 15) and assuming the same effect size distribution as in B and C, corresponding to the non-Lande case. Long-term stabilizing selection diminishes the contribution of alleles with lower initial MAF and amplifies that of alleles with higher initial MAF (red arrow); long-term, weak directional selection amplifies the contribution of alleles, regardless of their initial MAF (blue arrow). Consequently, the proportional long-term contribution of alleles with lower initial MAFs is diminished relative to their short-term contribution, an effect most pronounced for large effect sizes.

While our simple approximation underestimates the absolute contribution of alleles to phenotypic change, it does quite well at predicting the *relative* contribution of alleles with different effect sizes and initial frequencies (Fig. 5B and C). We can therefore generalize our approximations, such that

$$\Delta z_\infty(S, x_0) \approx C \cdot \frac{\partial \pi}{\partial x}(S, x_0) \cdot \Delta z_{t_1}(S, x_0) \approx 2C \cdot \Lambda f(S)/V_A(0) \qquad (15)$$

and

$$\Delta z_\infty(S) \approx C \cdot \Lambda f(S)/V_A(0), \qquad (16)$$

where the amplification factor C is determined by the requirement that the long-term change in mean phenotype equals the shift in optimum. This requirement yields

$$C \equiv V_A(0)/(2NU \int f(S) \cdot g(S) ds) = \int v(S) \cdot g(S) ds / \int f(S) \cdot g(S) ds.$$

Importantly, C depends only on the mutational distribution of effect sizes, $g(S)$. This generalized approximation follows from assuming that the amplification of the difference in the fixation probabilities of opposite alleles due to directional selection during the equilibration phase behaves as if *all* the frequency differences at the end of the rapid phase were amplified by a factor of C (SI Section 3.3). In that sense, C can be interpreted as an allelic measure of the deviation from Lande's approximation.

Our generalized approximation accounts for deviations from Lande's approximation that arise from large effect alleles but not for those that arise from large shifts in optimum (see section on "phenotypic response"). With larger shifts, the rapid phase would be longer and during this phase, directional selection would be stronger, leading to greater frequency differences between opposite alleles than our simple approximation suggests. Large frequency differences between opposite alleles also undermine the accuracy of our Taylor approximation for the difference in their fixation probabilities. In SI Section 3.3, we derive approximations that account for these effects, as well as those of weak directional selection during the equilibration phase, and are thus more accurate for any shift size and effect size distribution (Fig. 5B and S6). Nonetheless, our generalized simple approximation captures the salient features of the long-term allelic contribution to phenotypic adaptation (Fig. 5).

Notably, this approximation captures the dramatic turnover in the genetic basis of adaptation during the equilibration phase (Fig. 5B-D). Over the long-term, the short-term contribution of large effect alleles ($S \gtrsim 30$) is almost entirely wiped out, and is supplanted by the contribution of moderate effect alleles ($S \approx 5$) (Figs. 5B, C, and S7). Moreover, for any given effect size, the proportional long-term contribution of alleles with lower initial MAFs is diminished relative to their short-term contribution, all the more so for large effect sizes (Figs. 5D and S8). For instance, for an effect size $S = 35$, alleles with initial MAFs below 0.05 account for more than 99% of the short-term contribution but for only ~10% of the (much smaller) long-term contribution (Figs. 5D & S8).

We can understand this turnover by considering the effects of stabilizing selection during the equilibration phase. As noted, stabilizing selection induces selection against minor alleles, which weakens as MAF increase and becomes 0 at MAF = ½. Now consider how it affects a pair of alleles with opposite, moderate or large effect. If their initial frequencies are very low, both alleles will have low MAFs at the end of the rapid phase. Consequently, they will both be strongly selected against during the equilibration phase and will almost certainly go extinct. In the long-run, their expected contribution to phenotypic adaptation is therefore diminished. In contrast, if the alleles' initial MAF is sufficiently high, the relative increase in the aligned allele's frequency by the end of the rapid phase causes it to be subject to substantially weaker selection than is the opposing allele. In the extreme in which the aligned allele has exceeded frequency ½, the direction of selection on it is even reversed. In such cases, the pair's expected contribution to phenotypic adaptation will be amplified. This reasoning suggests that, for a given effect size, there is a critical initial MAF < ½ such that the long-term contribution of alleles that start above it is amplified and the contribution of those that start below it is diminished.

Weak directional selection during the equilibration phase also affects this behavior, however, by amplifying the long-term contribution of alleles with any initial MAF (Fig. 5D). Consequently, for a given amplification factor $C > 1$, the contribution of alleles with sufficiently small effect sizes are amplified for any initial MAF, as $\Delta z_{\infty}(S, x_0) > \Delta z_{t_1}(S, x_0)$ for any $x_0 > 0$ (Fig. 5D). In turn, for sufficiently large effect sizes, the curves for $\Delta z_{\infty}(S, x_0)$ and $\Delta z_{t_1}(S, x_0)$ intersect (Fig. 5D) and thus the critical frequency exists; in our approximation, it is the frequency where $C \cdot \partial\pi/\partial x\,(S, x_0) = 1$ (Eq. 15 and Fig. 5D). These considerations explain why, for sufficiently large effect sizes, the long-term contribution of alleles with low initial MAFs is diminished relative to their short-term contribution.

The turnover among alleles with different effect sizes can be explained in similar terms. Alleles with large effect sizes almost always start from low MAF, because they are subject to strong stabilizing selection before the shift (Figs. 4B). They are therefore highly unlikely to exceed the critical initial frequency, and their expected long-term contribution to

phenotypic adaptation is therefore diminished (Fig. 5A-C). However, changes to the frequencies of these alleles skew the phenotypic distribution, leading to prolonged, weak directional selection that amplifies the long-term contribution of small and moderate effect alleles (Fig. 5B and C). In our approximation, this amplification will occur for effect sizes that satisfy $C \cdot f(S) > v(S)$ (Eqs. 11 and 16). These considerations explain why the contributions of alleles with small and moderate effects supplant those of alleles with large effects over time (Fig. 5B and C). They also highlight that deviations from Lande's approximation are critical to understanding the allelic response, even when these deviations are modest and Lande's approximation provides a decent description of the phenotypic response.

Additional properties of the equilibration process. While the change in mean phenotype results from the preferential fixation of alleles whose effects are aligned with the shift, our results suggest that the expected number of these fixations is only slightly greater than the number of fixations with opposing effects (Fig. 6). The proportional excess of aligned fixations increases with the shift in optimum and allelic effect size, and for sufficiently large effect sizes, practically all fixations are of aligned alleles (Fig. 6C). Unless the contribution of alleles of small and moderate effects to genetic variance is negligible, however, the number of fixations of such large effect alleles and their contribution to phenotypic change will be small (Figs. 6A, B, and 5A-C). With the exception of this extreme case, we would therefore expect most fixations and contribution to phenotypic change to be of alleles with small and moderate effects (Figs. 6A, B, and 5A-C), for which the proportional excess of aligned fixations is modest (Fig. 6B).

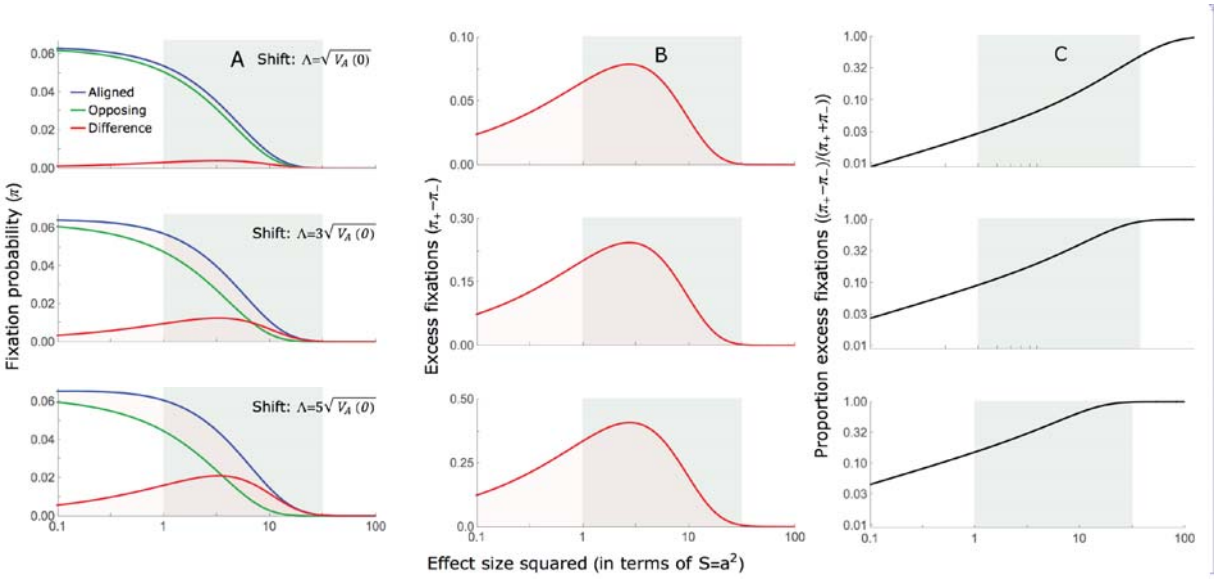

Figure 6. The excess of fixations of alleles whose effects are aligned with the shift, for three shifts in optimum. A) The fixation probabilities of aligned and opposing alleles as function of their effect size. For large effect sizes, the fixation probabilities become vanishingly small, whereas for small and moderate effect sizes, the difference in the fixation probabilities of alleles with opposing effects is small. B) Close up on the difference in the fixation probabilities between alleles with opposing effects. C) The proportional excess of fixations of aligned alleles among all fixations as a function of effect size.

Over the long-run, these fixations move the mean phenotype all the way to the new optimum, and genetic variation around this optimum returns to its steady state form. The time scale for equilibration is set by the fixation time of the alleles that underlie long-term phenotypic change. We can use the diffusion approximation to calculate the expected time to fixation of alleles with a given effect size and direction (SI Section 3.3; Fig. 7). As the perturbations to allele frequencies due to directional selection are fairly minor, the expected time to fixation of alleles with a given effect size is well approximated by their fixation time at steady state, i.e., assuming that their initial frequencies follow the equilibrium distribution and that they are only subject to stabilizing selection afterwards (Fig. 7). Because selection accelerates fixation (64), fixation times are longer for alleles with smaller effect sizes. We can therefore think about the equilibration process as taking

different times for alleles with distinct effect sizes. Based on the expected fixation times for alleles with small effects, which are the longest, we would expect steady state genetic variation around the new optimum to be restored by on the order of $\sim 4N$ generations after the shift in optimum (Fig. 7).

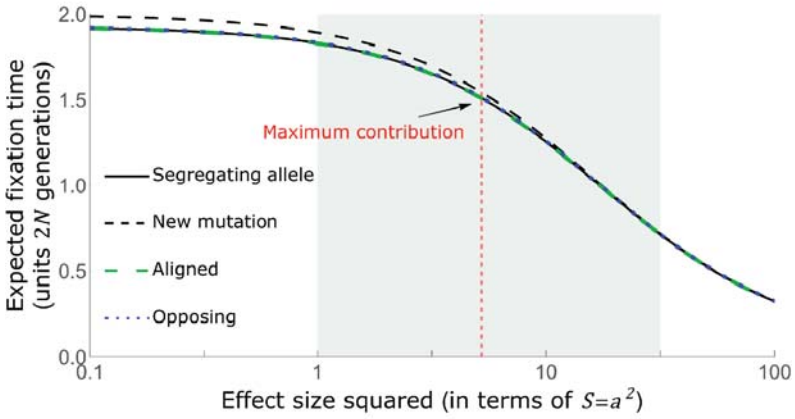

Figure 7. The equilibration time of alleles with different effect sizes. The effects of directional selection on opposite alleles are shown assuming Lande's approximation with a shift of $\Lambda = \sqrt{V_A(0)} = 25 \cdot \delta$. In Fig. S9 we show these effects for other parameter choices, including cases in which the phenotypic dynamic deviates from Lande's approximation. The frequency increase of aligned relative to opposing alleles shortens the fixation times of the aligned alleles and lengthens those of the opposing ones. However, these effects are minor (too minor to be seen in this case) and the fixation times of both aligned and opposing alleles are well approximated by the fixation times at steady state. The fixation times are slightly shorter than those for newly arising mutations because most of the alleles that fix are already segregating when the shift in optimum occurs.

Discussion

Here, we investigated the phenotypic and genetic adaptive response to selection on a highly polygenic quantitative trait in a simple yet highly relevant setting, in which a sudden change in environment shifts the trait's optimal value. The phenotypic response to selection was previously studied by Lande (35). He predicted that after the shift the population's mean phenotype will approach the new optimum exponentially, at a rate that is proportional to the additive genetic variance in the trait. This prediction, however, is predicated on the idealized infinitesimal model, in which the genetic variance arises from

an infinite number of segregating loci with infinitesimally small effects. We have extended this work to more realistic assumptions of a large but finite number of loci with a distribution of finite effect sizes. We found that Lande's prediction works well when the shift in optimum is not too large relative to genetic variance in the trait and the variance is dominated by loci with small and moderate effect sizes, which are defined based on the selection acting on these loci before the shift. When the shift is greater or when loci with large effects contribute markedly to genetic variance, the initial, rapid change in mean phenotype is followed by a pronounced quasi-static phase, governed by changes to the 3rd moment of the phenotypic distribution, in which the mean phenotype takes much longer to catch up to the new optimum.

We also characterized the genetic basis of these adaptive phenotypic changes. The closest previous work assumed an infinite population size (36, 37, 40, 43), and we found that relaxing this assumption leads to entirely different behavior. Notably, in infinite populations, small effect alleles whose frequencies before the shift are dominated by mutation and equal ½ make the greatest contribution to phenotypic change after the shift (see Introduction). In contrast, in any real (finite) population, the frequencies of such small effect alleles are dominated by genetic drift rather than mutation. More generally, variation in allele frequencies due to genetic drift, which is absent in infinite populations, critically affects the allelic response to selection.

To study the allelic response, we divided it into two periods: a rapid phase, immediately after the shift, and a subsequent, prolonged equilibration phase. During the rapid phase, the population's mean distance to the optimum is substantial and changes rapidly. Directional selection on the trait increases the frequency of minor alleles whose effects are aligned with the shift relative to minor alleles with opposing effects (given the same effect size and initial frequency). By the end of the rapid phase, the cumulative effect of these frequency differences pushes the mean phenotype close to the new optimum, but because this effect is spread over myriad alleles, the frequency difference between any individual pair of opposing alleles is fairly small. Specifically, we found that an allele's contribution to phenotypic change is proportional to its contribution to phenotypic variance before the

shift, implying that alleles with moderate and large effect sizes make the greatest per site contributions to phenotypic change, while alleles with moderate effect sizes experience the greatest frequency changes. The expected frequency differences between opposing alleles is amplified by prolonged, weak directional selection during the subsequent equilibration phase, and this amplification is pronounced when the phenotypic approach to the new optimum deviates markedly from Lande's approximation.

Over the long run, stabilizing selection on the trait transforms these small frequency differences into a small excess of fixed aligned alleles relative to opposing ones, and cumulatively this excess moves the population mean all the way to the new optimum. This transformation process involves a massive turnover in the properties of the contributing alleles. Notably, the transient contributions of large effect alleles are supplanted by contributions of fixed moderate, and to a lesser extent, small effect alleles. This process takes on the order of $4N_e$ generations, after which the steady state architecture of genetic variation around the new optimum is restored.

Our finding that large effect alleles almost never sweep to fixation appears at odds with the results of previous studies of similar models. These discrepancies are largely explained by earlier papers considering settings that violate our assumptions, notably about evolutionary parameter ranges. For instance, some studies assume that large effect alleles segregate at high frequencies before the shift in optimum (e.g., (65)), which is presumably uncommon in natural populations and in any case, violates our assumption that the trait is at steady state before the shift. Other models implicitly consider quantitative traits of intermediate genetic complexity; while such traits likely exist, there are to our knowledge few well-established examples. Notably, Thornton (38) observes sweeps in cases in which the trait is not highly polygenic (violating our assumption that $\sqrt{2NU} \gg 1$). Relatedly, Chevin and Hospital (66) observe sweeps in cases in which a single newly arising mutation of large effect contributes substantially to genetic variance, which violates our assumptions that genetic variation is highly polygenic and is not predominantly effectively neutral (i.e., that alleles with $S \gtrsim 1$ contribute substantially). Although it remains to be seen, we believe

that this architecture is much less common, given mounting evidence, reviewed in the Introduction, which suggests that traits are often highly polygenic, and other considerations, notably estimates of persistence time (3, 4) and inferences based on human GWASs (49, 54), which indicate that quantitative genetic variation is not predominantly neutral.

Lastly, Stetter et al. (39) considered a huge shift in the optimal trait value (e.g., of ~90 phenotypic standard deviations), resulting in a massive drop in fitness (violating our assumption that $\Lambda < \sqrt{V_S}$)—although shifts in optimum need not be that large to result in some fixations of large effect alleles. While there are many examples of rapid and large environmental fluctuations, e.g., due to seasonal fluctuations or weather systems, they occur on a much shorter time scale than fixation (although they might have some effect on genetic architecture; see below). In turn, little is known about the magnitude of shifts in optimal trait values over the time scales of large effect, beneficial fixations. While it seems plausible that moderate shifts, which fall within our assumed parameter ranges, are common, we cannot rule out that larger shifts are common as well. The response to such larger shifts is not covered by our analysis and clearly warrants further study.

Other factors that we have not considered may also affect polygenic adaptation. Most notable among them is pleiotropy. Given that quantitative genetic variation affecting one trait often affects many other traits (3, 19, 67-69), alleles that would have been positively selected because of their effect on the trait under directional selection may be selected against because of their adverse effects on other traits. Moreover, pleiotropy is known to affect the genetic architecture of a given trait at steady state (49), which we have shown to shape the allelic response to selection on that trait. Pleiotropy is therefore likely to affect which alleles contribute to phenotypic change at the different phases of polygenic adaptation (see (70) for related considerations for simple traits). Linkage disequilibrium (LD) may have an effect as well, perhaps most notably for minor alleles with large effects, which start at low frequencies and experience strong directional selection during the rapid phase. Before the shift, large effect alleles located in genomic regions with low

recombination and high functional density are more likely to be in LD with, for example, alleles with countervailing effects on the focal trait (71) or deleterious effects on other traits. If this were the case, then directional selection during the rapid phase might be effectively weaker, because it would act on extended haplotypes rather than on individual alleles.

In addition, the demography of a population, notably its size, as well as the selection pressures on quantitative traits are likely to change over a shorter time scale than it takes the genetic architecture of complex traits to equilibrate. When these changes occur over the $\sim 4N_e$ generations preceding a shift in optimal trait value, they would affect the genetic architecture of the trait and consequently its response to selection. Changes in population size influence the number of segregating sites affecting a trait and the distribution of their frequencies and contributions to variance, with more recent population sizes affecting strongly selected variation more than weakly selected variation (3, 72-74). The effects of varying selection will depend on the attributes of this variation in ways that await further study.

While the effects of all of these factors on the response to a shift in optimum warrant investigation, we expect the response to follow from the principles we outlined. Notably, we expect the short-term contribution of alleles to phenotypic change to be proportional to their contribution to variance before the shift, and their long-term contribution to arise from differences between the fixation probabilities of alleles with opposite effects, caused by the opposing effects of directional selection on their frequencies. Thus, while all of these factors are likely to affect the response, we expect the main features of the dynamics we portrayed to remain largely intact. These features include the role of the 3rd moment of the phenotypic distribution in slowing down phenotypic adaptation near the new optimum; the transient contribution of large effect alleles to phenotypic adaptation; and the long-term importance of alleles with moderate effects.

As polygenic adaptation in quantitative traits is likely ubiquitous, our conclusions have potentially important implications. One is that, contrary to adaptation mediated by

selective sweeps of initially rare, large effect, beneficial alleles (10, 11, 75-78), polygenic adaptation might have minor effects on patterns of neutral diversity. The effects of selected alleles on neutral diversity at linked loci follow from their trajectories (79). Our results indicate that directional selection on a highly polygenic trait introduces only small changes to allele frequencies at individual loci, which amount to minor perturbations to the allelic trajectories expected under stabilizing selection at steady state (also see (38, 66)). Indeed, alleles with large effects exhibit only small, transient changes; among alleles with moderate effects, there is a long-term excess of fixations of those whose effects are aligned with the shift relative to those whose effects are opposed, but this excess is small (Fig. 6) and the trajectories of the alleles that fix are largely driven by weak stabilizing selection (Fig. 7). Thus, our results indicate that the effects of polygenic adaptation on neutral diversity should be minor (other than perhaps for massive shifts in optimal trait values, as noted above).

In contrast, long-term stabilizing selection on quantitative traits likely has substantial effects on neutral diversity patterns. Specifically, selection against minor alleles induced by stabilizing selection may well be a major source of background selection and is expected to affect neutral diversity patterns in ways that are similar to those of background (purifying) selection from other selective origins (80-82).

Another implication of our results pertains to the search for the genetic basis of human adaptation, as well as adaptation in other species. Efforts to uncover the identity of individual adaptive genetic changes on the human lineage were guided by the notion that their identity would offer insight into what "made us human". Under the plausible assumption that many adaptive changes on the human lineage arose from selection on complex, quantitative traits, this approach may not be as informative as it appears (15, 19). Our results indicate that after a shift in the optimal trait value, the number of fixations of alleles whose effects are aligned to the shift are nearly equal to the number of alleles that are opposed (Fig. 6). Moreover, the alleles that fix are a largely random draw from the vastly greater number of alleles that affect the trait, both in the sense of being those that happened to segregate at high MAFs at the onset of selection and because of the

stochasticity of fixation. Thus, in this plausible scenario, it becomes meaningless to say that any given fixation was adaptive, and arguably uninteresting to focus on the particular subset of alleles that happened to reach fixation. In contrast, identifying the *traits* that experienced adaptive changes promises to provide important insights. Recent efforts to do so pool the signatures of frequency changes over many loci that were found to be associated with a given trait in GWAS (20-26), an exciting approach that has also proven to be technically challenging (27, 28). A better understanding of the process of polygenic adaptation should help to guide such efforts.

Acknowledgements. We thank Guy Amster, Jeremy Berg, Nick Barton, Yuval Simons and Molly Przeworski for many helpful discussions, and Jeremy Berg, Joachim Hermisson, Will Milligan, Yuval Simons, Leo Speidel and Molly Przeworski for comments on the manuscript. This work was funded by NIH grant GM115889 to GS.

References

1. Fisher RA (1930) *The genetical theory of natural selection* (The Clarendon press, Oxford,) pp xiv, 272 p.
2. Wright S (1931) Evolution in Mendelian Populations. *Genetics* 16(2):97-159.
3. Sella G & Barton NH (2019) Thinking About the Evolution of Complex Traits in the Era of Genome-Wide Association Studies. *Annual review of genomics and human genetics*.
4. Walsh B & Lynch M (2018) *Evolution and Selection of Quantitative Traits* (Oxford University Presss).
5. Weber KE & Diggins LT (1990) Increased selection response in larger populations. II. Selection for ethanol vapor resistance in Drosophila melanogaster at two population sizes. *Genetics* 125(3):585-597.
6. Barton NH & Keightley PD (2002) Understanding quantitative genetic variation. *Nature reviews. Genetics* 3(1):11-21.
7. Hill WG (2016) Is continued genetic improvement of livestock sustainable? *Genetics* 202(3):877-881.
8. Hill WG & Kirkpatrick M (2010) What animal breeding has taught us about evolution. *Annu Rev Ecol Evol S* 41:1-19.
9. Zhang XS & Hill WG (2005) Predictions of patterns of response to artificial selection in lines derived from natural populations. *Genetics* 169(1):411-425.
10. Maynard Smith JM & Haigh J (1974) The hitch-hiking effect of a favourable gene. *Genetical research* 23(1):23-35.

11. Kaplan NL, Hudson RR, & Langley CH (1989) The "hitchhiking effect" revisited. *Genetics* 123(4):887-899.
12. Voight BF, Kudaravalli S, Wen X, & Pritchard JK (2006) A map of recent positive selection in the human genome. *PLoS biology* 4(3):e72.
13. Coop G, *et al.* (2009) The role of geography in human adaptation. *PLoS Genet* 5(6):e1000500.
14. Hernandez RD, *et al.* (2011) Classic selective sweeps were rare in recent human evolution. *Science* 331(6019):920-924.
15. Pritchard JK, Pickrell JK, & Coop G (2010) The genetics of human adaptation: hard sweeps, soft sweeps, and polygenic adaptation. *Curr Biol* 20(4):R208-215.
16. Pritchard JK & Di Rienzo A (2010) Adaptation - not by sweeps alone. *Nature reviews. Genetics* 11(10):665-667.
17. Shi H, Kichaev G, & Pasaniuc B (2016) Contrasting the Genetic Architecture of 30 Complex Traits from Summary Association Data. *American journal of human genetics* 99(1):139-153.
18. Loh PR, *et al.* (2015) Contrasting genetic architectures of schizophrenia and other complex diseases using fast variance-components analysis. *Nature genetics* 47(12):1385-1392.
19. Boyle EA, Li YI, & Pritchard JK (2017) An Expanded View of Complex Traits: From Polygenic to Omnigenic. *Cell* 169(7):1177-1186.
20. Turchin MC, *et al.* (2012) Evidence of widespread selection on standing variation in Europe at height-associated SNPs. *Nature genetics* 44(9):1015-1019.
21. Berg JJ & Coop G (2014) A population genetic signal of polygenic adaptation. *PLoS Genet* 10(8):e1004412.
22. Robinson MR, *et al.* (2015) Population genetic differentiation of height and body mass index across Europe. *Nature genetics* 47(11):1357-1362.
23. Field Y, *et al.* (2016) Detection of human adaptation during the past 2000 years. *Science* 354(6313):760-764.
24. Edge MD & Coop G (2019) Reconstructing the History of Polygenic Scores Using Coalescent Trees. *Genetics* 211(1):235-262.
25. Speidel L, Forest M, Shi S, & Myers S (2019) A method for genome-wide genealogy estimation for thousands of samples. (bioRxiv).
26. Berg JJ, Zhang X, & Coop G (2017) Polygenic Adaptation has Impacted Multiple Anthropometric Traits. (bioRxiv).
27. Berg JJ, *et al.* (2019) Reduced signal for polygenic adaptation of height in UK Biobank. *eLife* 8:e39725.
28. Sohail M, *et al.* (2019) Polygenic adaptation on height is overestimated due to uncorrected stratification in genome-wide association studies. *eLife* 8.
29. Robertson A (1960) A theory of limits in artificial selection. *P Roy Soc B-Biol Sci* 153:234-249.
30. Burger R (1999) Evolution of genetic variability and the advantage of sex and recombination in changing environments. *Genetics* 153(2):1055-1069.
31. Kopp M & Hermisson J (2009) The genetic basis of phenotypic adaptation I: fixation of beneficial mutations in the moving optimum model. *Genetics* 182(1):233-249.

32. Matuszewski S, Hermisson J, & Kopp M (2015) Catch Me if You Can: Adaptation from Standing Genetic Variation to a Moving Phenotypic Optimum. *Genetics* 200(4):1255-1274.
33. Burger R & Lynch M (1995) EVOLUTION AND EXTINCTION IN A CHANGING ENVIRONMENT: A QUANTITATIVE-GENETIC ANALYSIS. *Evolution; international journal of organic evolution* 49(1):151-163.
34. Jain K & Devi A (2018) Polygenic adaptation in changing environments. *Europhysics Letters* 123(4).
35. Lande R (1976) Natural-Selection and Random Genetic Drift in Phenotypic Evolution. *Evolution; international journal of organic evolution* 30(2):314-334.
36. de Vladar HP & Barton N (2014) Stability and response of polygenic traits to stabilizing selection and mutation. *Genetics* 197(2):749-767.
37. Jain K & Stephan W (2017) Rapid adaptation of a polygenic trait after a sudden environmental shift. *Genetics* 206(1):389-406.
38. Thornton KR (2018) Polygenic adaptation to an environmental shift: temporal dynamics of variation under Gaussian stabilizing selection and additive effects on a single trait. (bioRxiv).
39. Stetter MG, Thornton K, & Ross-Ibarra J (2018) Genetic architecture and selective sweeps after polygenic adaptation to distant trait optima. *PLoS Genet* 14(11):e1007794.
40. Jain K & Stephan W (2015) Response of Polygenic Traits Under Stabilizing Selection and Mutation When Loci Have Unequal Effects. *G3 (Bethesda, Md.)* 5(6):1065-1074.
41. Barton NH & de Vladar HP (2009) Statistical mechanics and the evolution of polygenic quantitative traits. *Genetics* 181(3):997-1011.
42. Bod'ova K, Tkacik G, & Barton NH (2016) A General Approximation for the Dynamics of Quantitative Traits. *Genetics* 202(4):1523-1548.
43. Jain K & Stephan W (2017) Modes of Rapid Polygenic Adaptation. *Mol Biol Evol* 34(12):3169-3175.
44. Wright S (1935) The analysis of variance and the correlations between relatives with respect to deviations from an optimum. *Journal of Genetics* 30(2):243-256.
45. Robertson A (1956) The effect of selection against extreme deviants based on deviation or on homozygosis. *Journal of Genetics* 54:236-248.
46. Turelli M (1984) Heritable genetic variation via mutation-selection balance: Lerch's zeta meets the abdominal bristle. *Theor Popul Biol* 25(2):138-193.
47. Keightley PD & Hill WG (1988) Quantitative genetic variability maintained by mutation-stabilizing selection balance in finite populations. *Genetical research* 52(1):33-43.
48. Johnson T & Barton N (2005) Theoretical models of selection and mutation on quantitative traits. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences* 360(1459):1411-1425.
49. Simons YB, Bullaughey K, Hudson RR, & Sella G (2018) A population genetic interpretation of GWAS findings for human quantitative traits. *PLoS biology* 16(3):e2002985.
50. Falconer DS & Mackay TFC (1996) *Introduction to quantitative genetics* (Essex, England: Longman).

51. Lynch M & Walsh B (1998) *Genetics and analysis of quantitative traits* (Sinauer, Sunderland, Mass.) pp xvi, 980 p.
52. Burger R (2000) *The mathematical theory of selection, recombination, and mutation* (Wiley).
53. Wright S (1935) Evolution in populations in approximate equilibrium. *Journal of Genetics* 30:257-266.
54. Zeng J, *et al.* (2018) Signatures of negative selection in the genetic architecture of human complex traits. *Nature genetics* 50(5):746-753.
55. Barton N & Turelli M (1987) Adaptive landscapes, genetic distance and the evolution of quantitative characters. *Genetics Research* 49(2):157-173.
56. Burger R (1991) Moments, cumulants, and polygenic dynamics. *Journal of mathematical biology* 30(2):199-213.
57. Fisher RA (1918) The Correlation Between Relatives on the Supposition of Mendelian Inheritance. *Transactions of the Royal Society of Edinburgh* 52:399-433.
58. Bulmer MG (1985) *The mathematical theory of quantitative genetics* (Oxford University Press, Oxford).
59. Barton NH, Etheridge AM, & Veber A (2017) The infinitesimal model: Definition, derivation, and implications. *Theor Popul Biol* 118:50-73.
60. Crow JF & Kimura M (1970) *An introduction to population genetics theory* (Harper & Row, New York,) pp xiv, 591 p.
61. Ewens WJ (2004) *Mathematical population genetics* (Springer, New York) 2nd Ed pp v. <1- >.
62. Barton N (1986) The maintenance of polygenic variation through a balance between mutation and stabilizing selection. *Genetics Research* 47(3):209-216.
63. Charlesworth B (2013) Stabilizing selection, purifying selection, and mutational bias in finite populations. *Genetics* 194(4):955-971.
64. Maruyama T (1974) The age of a rare mutant gene in a large population. *American journal of human genetics* 26(6):669-673.
65. Christodoulaki E, Barghi N, & Schloetterer C (2019) Distance to trait optimum is a crucial factor determining the genomic signature of polygenic adaptation. (BioRxiv).
66. Chevin LM & Hospital F (2008) Selective sweep at a quantitative trait locus in the presence of background genetic variation. *Genetics* 180(3):1645-1660.
67. Bulik-Sullivan B, *et al.* (2015) An atlas of genetic correlations across human diseases and traits. *Nature genetics* 47(11):1236-1241.
68. Pickrell JK, *et al.* (2016) Detection and interpretation of shared genetic influences on 42 human traits. *Nature genetics* 48(7):709-717.
69. Liu X, Li YI, & Pritchard JK (2019) Trans Effects on Gene Expression Can Drive Omnigenic Inheritance. *Cell* 177(4):1022-1034.e1026.
70. Otto SP (2004) Two steps forward, one step back: the pleiotropic effects of favoured alleles. *Proc Biol Sci* 271(1540):705-714.
71. Lande R (1975) The maintenance of genetic variability by mutation in a polygenic character with linked loci. *Genetical research* 26(3):221-235.
72. Simons YB, Turchin MC, Pritchard JK, & Sella G (2014) The deleterious mutation load is insensitive to recent population history. *Nature genetics* 46(3):220-224.

73. Simons YB & Sella G (2016) The impact of recent population history on the deleterious mutation load in humans and close evolutionary relatives. *Current opinion in genetics & development* 41:150-158.
74. Lohmueller KE (2014) The impact of population demography and selection on the genetic architecture of complex traits. *PLoS Genet* 10(5):e1004379.
75. Braverman JM, Hudson RR, Kaplan NL, Langley CH, & Stephan W (1995) The hitchhiking effect on the site frequency spectrum of DNA polymorphisms. *Genetics* 140(2):783-796.
76. Hermisson J & Pennings PS (2005) Soft sweeps: molecular population genetics of adaptation from standing genetic variation. *Genetics* 169(4):2335-2352.
77. Coop G & Ralph P (2012) Patterns of neutral diversity under general models of selective sweeps. *Genetics* 192(1):205-224.
78. Berg JJ & Coop G (2015) A Coalescent Model for a Sweep of a Unique Standing Variant. *Genetics* 201(2):707-725.
79. Barton NH (2000) Genetic hitchhiking. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences* 355(1403):1553-1562.
80. Charlesworth B, Morgan MT, & Charlesworth D (1993) The effect of deleterious mutations on neutral molecular variation. *Genetics* 134(4):1289-1303.
81. Hudson RR & Kaplan NL (1995) Deleterious background selection with recombination. *Genetics* 141(4):1605-1617.
82. McVean GA & Charlesworth B (2000) The effects of Hill-Robertson interference between weakly selected mutations on patterns of molecular evolution and variation. *Genetics* 155(2):929-944.

Symbol	Definition
N	Population size
U	Expected number of mutations per gamete per generation affecting the trait
V_S	Width of the Gaussian fitness function
δ	Magnitude of fluctuations around the optimum at steady state ($=\sqrt{V_S}/(2N)$)
a	An allele's effect on the trait
S	An allele's scaled selection coefficient at steady-state ($=a^2$ in units of δ^2)
$g(S)$	The mutational distribution of scaled selection coefficients
Λ	The size of the shift in optimum
t	Time after shift in optimum
$V_A(t)$	The additive genetic variance (and second moment of the trait distribution)
$\mu_3(t)$	The third moment of the trait distribution
$D(t)$	Distance between the mean phenotype and optimum
$D_L(t)$	Lande's approximation for $D(t)$
t_1	The time at the end of the rapid phase
$x_t^+(S,x_0)$ $x_t^-(S,x_0)$	Expected frequency of an allele that is aligned (+) or opposed (-) to the shift with selection coefficient S and initial frequency x_0
$\Delta x_t^*(S,x_0)$	Expected frequency difference between opposite alleles ($=x_t^+(S,x_0)-x_t^-(S,x_0)$)
$\Delta z_t^*(S,x_0)$	Expected contribution to phenotypic change of a pair of opposite alleles
$v^*(S,x)$	An allele's contribution to phenotypic variance ($=2a^2x(1-x)$)
$\Delta z_t(S)$ $\Delta z_t(S,x_0)$	Expected contribution to phenotypic change per unit mutational input of alleles with selection coefficient S, or with selection coefficient S and initial frequency x_0
$v(S)$ $v(S,x_0)$	Steady state density of phenotypic variance per unit mutational input of alleles with selection coefficient S, or with selection coefficient S and frequency x_0
$\pi(S,x)$	Fixation probability under stabilizing selection with selection coefficient S and initial frequency x
$f(S)$	Relative long-term contribution to phenotypic change
C	Amplification factor of the long-term contribution to phenotypic change ($=\int_S v(S)g(S)\,dS/\int_S f(S)g(S)\,dS$)

Table S1. Summary of notation.